Image Data Compression for Covariance and Histogram Descriptors

Matt J. Kusner, Nicholas I. Kolkin
Washington University in St. Louis
{mkusner,n.kolkin}@wustl.edu

Stephen Tyree
NVIDIA Research
styree@nvidia.com

Kilian Q. Weinberger
Washington University in St. Louis
kilian@wustl.edu

arXiv:1412.1740v2 [stat.ML] 23 May 2015

Abstract

Covariance and histogram image descriptors provide an effective way to capture information about images. Both excel when used in combination with special purpose distance metrics. For covariance descriptors these metrics measure the distance along the non-Euclidean Riemannian manifold of symmetric positive definite matrices. For histogram descriptors the Earth Mover's distance measures the optimal transport between two histograms. Although more precise, these distance metrics are very expensive to compute, making them impractical in many applications, even for data sets of only a few thousand examples. In this paper we present two methods to compress the size of covariance and histogram datasets with only marginal increases in test error for k-nearest neighbor classification. Specifically, we show that we can reduce data sets to $\mathbf{16\%}$ and in some cases as little as $\mathbf{2\%}$ of their original size, while approximately matching the test error of kNN classification on the full training set. In fact, because the compressed set is learned in a supervised fashion, it sometimes even outperforms the full data set, while requiring only a fraction of the space and drastically reducing test-time computation.

1. Introduction

In the absence of sufficient data to learn image descriptors directly, two of the most influential classes of feature descriptors are (i) the histogram and (ii) the covariance descriptor. Histogram descriptors are ubiquitous in computer vision [2, 10, 27, 32, 33]. These descriptors may be designed to capture the distribution of image gradients throughout an image, or they may result from a visual bag-of-words representation [11, 12, 29]. Covariance descriptors, and more generally, symmetric positive definite (SPD) matrices, are often used to describe structure tensors [15], diffusion tensors [37] or region covariances [45]. The latter are particularly well suited for the task of object detection from a variety of viewpoints and illuminations.

In this paper, we focus on the k-nearest neighbor (kNN) classifier with histogram or covariance image descriptors. Computing a nearest neighbor or simply comparing a pair of histograms or SPD matrices is non-trivial. For histogram descriptors, certain bins may be individually similar/dissimilar to other bins. Therefore, the Euclidean distance is often a poor measure of distance as it cannot measure such bin-wise dissimilarity. In the case of covariance descriptors, SPD matrices lie on a convex half-cone—a non-Euclidean Riemannian manifold embedded inside a Euclidean space. Measuring distances between SPD matrices with the straight-forward Euclidean metric ignores the underlying manifold structure of the data and tends to systematically under-perform in classification tasks [47].

Histogram and covariance descriptors excel if their underlying structure is incorporated into the distance metric. Recently, there have been a number of proposed histogram distances [34, 36, 38, 43]. Although these yield strong improvements in kNN classification accuracy (versus the Euclidean distance), these distances are often very costly to compute (*e.g.* super-cubic in the histogram dimensionality). Similarly, for SPD matrices there are specialized geodesic distances [6] and algorithms [22, 45] developed that operate on the SPD covariance manifold. Because of the SPD constraint, these methods often require significantly more time to make predictions on test data. This is especially true for kNN, as the computation of the geodesic distance along the SPD manifold requires an eigen-decomposition *for each individual pairwise distance*— a computation that needs to be repeated for all training inputs to classify a single test input.

Cherian et al. [6] improved the running time (in practice) of test classification by approximating the Riemmanian distance with a symmetrized log-determinant divergence. For low dimensional data, Bregman Ball Trees [4] can be adapted, however the performance deteriorates quickly as the dimensionality increases.

In this paper, we develop a novel technique to speed up k-nearest neighbor applications on covariance and histogram image features, one that can be used in concert with many other speedup methods. Our methods, called *Stochastic Covariance Compression* (SCC) and *Stochastic Histogram Compression* (SHC) learn a *compressed training set* of size m, such that $m \ll n$, which approximately

matches the performance of the kNN classifier on the original data. This new data set does not consist of original training samples; instead it contains new, artificially generated inputs which are explicitly designed for low kNN error on the training data. The original training set can be discarded after training and during test-time, as we only find the k-nearest neighbors among these artificial samples. This drastically reduces computation time and shrinks storage requirements. To facilitate learning compressed data sets we borrow the concept of stochastic neighborhoods, used in data visualization [21, 46] and metric learning [16], and leverage recent results from the machine learning community on data compression in Euclidean spaces [24].

We make three novel contributions: 1. we derive SCC and SHC, two new methods for compression of covariance and histogram data; 2. we devise efficient methods for solving the SCC and SHC optimizations using the Cholesky decomposition and a normalized change of variable; 3. we carefully evaluate both methods on several real world data sets and compare against an extensive set of state-of-the-art baselines. Our experiments show that SCC can often compress a covariance data set to about 16% of its original size, without increase in kNN test error. In some cases, SHC and SCC can match the kNN test error with only 2% of the training set size—leading to order-of-magnitude speedups during test time. Finally, because we learn the compressed set explicitly to minimize kNN error, in a few cases it even *outperforms* the full data set by achieving *lower test error*.

2. Covariance and Histogram Descriptors

We assume that we are given a set of d-dimensional feature vectors $\mathcal{F} = \{\mathbf{x}_1, \ldots, \mathbf{x}_{|\mathcal{F}|}\} \subset \mathbb{R}^d$ computed from a *single* input image. From these features, we compute covariance or histogram image descriptors.

Covariance descriptors represent $\mathcal{F}$ through the covariance matrix of the features,

$$\mathbf{X} = \frac{1}{|\mathcal{F}| - 1} \sum_{r=1}^{|\mathcal{F}|} (\mathbf{x}_r - \boldsymbol{\mu})(\mathbf{x}_r - \boldsymbol{\mu})^\top,$$

where $\boldsymbol{\mu} = \frac{1}{|\mathcal{F}|} \sum_{r=1}^{|\mathcal{F}|} \mathbf{x}_r$. For vectorial data, 'nearness' is often computed via the Euclidean distance or a learned Mahalanobis metric [48]. However, the Euclidean/Mahalanobis distance between two covariance matrices is a poor approximation to their true distance along the manifold of SPD matrices. A natural distance for covariance matrices is the Affine-Invariant Riemannian metric [6], a geodesic distance on the SPD manifold.

Definition 1. *Let $\mathcal{S}_+^d$ be the positive definite cone of matrices of rank d. The **Affine-Invariant Riemannian metric** (AIRM) between any two matrices $\mathbf{X}, \hat{\mathbf{X}} \in \mathcal{S}_+^d$ is $D_R(\mathbf{X}, \hat{\mathbf{X}}) = \| \log(\hat{\mathbf{X}}^{-1/2} \mathbf{X} \hat{\mathbf{X}}^{-1/2}) \|_F$.*

While the AIRM accurately describes the dissimilarity between two covariances along the SPD manifold, it requires an *eigenvalue decomposition for every input* $\hat{\mathbf{X}}$. The metric becomes intractable to compute even for moderately-sized covariance matrices (*e.g.,* computing $\hat{\mathbf{X}}^{-1/2} \in \mathcal{S}_+^d$ requires roughly $\mathrm{O}\big(d^3\big)$ time). To alleviate this computational burden, a distance metric with similar theoretical properties has been proposed by [6], called the Jensen-Bregman LogDet Divergence (JBLD),

$$D_J(\mathbf{X}, \hat{\mathbf{X}}) = \log \Big| \frac{\mathbf{X} + \hat{\mathbf{X}}}{2} \Big| - \frac{1}{2} \log |\mathbf{X}\hat{\mathbf{X}}|. \qquad (1)$$

Cherian et al. [6] demonstrate that for nearest neighbor classification, using JBLD as a distance has performance nearly identical to the AIRM but is much faster in practice and asymptotically requires $O(d^{2.37})$ computation [7].

Histogram descriptors are a popular alternative to covariance representations. Assume we again have a set of d-dimensional features for an image $\mathcal{F} = \{\mathbf{x}_1, \ldots, \mathbf{x}_{|\mathcal{F}|}\}$. Further, let the collection of all such features for all n images in a training set be $\boldsymbol{\mathcal{F}} = \mathcal{F}_i \cup \ldots \cup \mathcal{F}_n$ (where $\mathcal{F}_i$ are the features for image i). To construct the visual bag-of-words representation we cluster all features in $\boldsymbol{\mathcal{F}}$ into K centroids $\mathbf{c}_1, \ldots, \mathbf{c}_K$ (e.g., via k-means), where these centroids are often referred to as a *codebook* [12]. Using this codebook the visual bag-of-words representation $\mathbf{h}_i$ of an image i is a K-dimensional vector, where element h_{ij} is a count of how many features in the bag $\mathcal{F}_i$ have $\mathbf{c}_j$ as the nearest centroid.

Arguably one of the most successful histogram distances is the Earth Mover's Distance (EMD) [39], which has been used to achieve impressive results for image classification and retrieval [30, 31, 39, 35]. EMD constructs a distance between two histograms by 'lifting' a bin-to-bin distance, called the *ground distance* $\mathbf{M}$, where $M_{ij} \geq 0$, to a full histogram distance. Specifically, for two histogram vectors $\mathbf{h}$ and $\mathbf{h}'$ the EMD distance is the solution to the following linear program:

$$\min_{\mathbf{T} \geq 0} \ \mathrm{tr}(\mathbf{TM}) \ \text{ s.t. } \mathbf{T1} = \mathbf{h} \text{ and } \mathbf{T}^\top \mathbf{1} = \mathbf{h}', \qquad (2)$$

where $\mathbf{T}$ is the transportation matrix and $\mathbf{1}$ is a vector of ones. Each element T_{ij} describes the amount of mass moved from h_i to h'_j, for the vectors to match exactly. One example ground distance for the visual bag-of-words representation is the Euclidean distance between the centroid vectors $M_{ij} \triangleq \|\mathbf{c}_i - \mathbf{c}_j\|_2$. When the ground distance is a metric, it can be shown that the EMD is also a metric [39].

In practice, one limitation of the EMD distance is its high computational complexity. Cuturi et al. [8] therefore introduce the Sinkhorn Distance, which involves a regularized version of the EMD optimization problem:

$$\min_{\mathbf{T} \geq 0} \mathrm{tr}(\mathbf{TM}) - \frac{1}{\lambda} h(\mathbf{T}), \text{ s.t. } \mathbf{T1} = \mathbf{h} \text{ and } \mathbf{T}^\top \mathbf{1} = \mathbf{h}' \quad (3)$$

where $h(\mathbf{T}) = -\text{tr}(\mathbf{T}\log(\mathbf{T}))$ is the entropy of the transport $\mathbf{T}$. The Sinkhorn distance between $\mathbf{h}$ and $\mathbf{h}'$ is $D_S(\mathbf{h}, \mathbf{h}') = \text{tr}(\mathbf{T}^\lambda \mathbf{M})$, where $\mathbf{T}^\lambda$ is the solution to (3). The solution is an arbitrarily close upper bound to the exact EMD solution (by increasing λ) and the optimization problem is shown to be at least an order of magnitude faster to compute than the EMD linear program (2). Specifically, Cuturi et al. introduce a simple iterative algorithm to solve eq.(3) in time $\mathrm{O}\big(d^2 i\big)$, where d is the size of the histograms and i is the number of iterations of the algorithm. This is compared to $\mathrm{O}\big(d^3 \log d\big)$ complexity of the EMD optimization problem [35]. In practice, each algorithm iteration is a matrix scaling computation that can be performed between multiple histograms simultaneously. This means that the algorithm is parallel and can be efficiently computed on modern hardware architectures (*i.e.*, multi-core CPUs and GPUs) [8].

3. Covariance compression

In this section we detail our covariance compression technique: *Stochastic Covariance Compression* (SCC). SCC uses a stochastic neighborhood to compress the training set from n input covariances to m 'compressed' covariances. After learning, the original training set can be discarded and all future classifications are made just using the compressed inputs. Since $m \ll n$, the complexity of test-time classification is drastically reduced, from $\mathrm{O}\big(nd^{2.37}\big)$ to $\mathrm{O}\big(md^{2.37}\big)$, where $\mathrm{O}\big(d^{2.37}\big)$ is the asymptotic complexity of computing a single JBLD distance in eq. (1).

Assume we are given a training set of n covariance matrices $\{\mathbf{X}_1, \ldots, \mathbf{X}_n\} \subset \mathbb{R}^{d\times d}$ with corresponding labels $y_1, \ldots, y_n$. Our goal is to learn a compressed set of m covariance matrices $\{\hat{\mathbf{X}}_1, \ldots, \hat{\mathbf{X}}_m\} \subset \mathbb{R}^{d\times d}$ with labels $\hat{y}_1, \ldots, \hat{y}_m$. To initialize $\hat{\mathbf{X}}_j$, we randomly sample m covariance matrices from our training data set and copy their associated labels for each $\hat{y}$. We optimize these synthetic inputs $\hat{\mathbf{X}}_j$ to minimize the kNN classification error. The kNN classification error is non-continuous and non-differentiable with respect to $\hat{\mathbf{X}}_j$, but we can introduce a stochastic neighborhood, as proposed by Hinton and Roweis [21], to "soften" the neighborhood assignment and allow optimization on kNN error. Specifically, we place a radial basis function around each input $\mathbf{X}_i$ and proceed as if the nearest prototypes $\hat{\mathbf{X}}_j$ are assigned randomly. For a given $\mathbf{X}_i$, the probability that $\hat{\mathbf{X}}_j$ is picked as the nearest neighbor is denoted

$$p_{ij} = \frac{e^{-\gamma^2 D_J(\mathbf{X}_i, \hat{\mathbf{X}}_j)}}{\sum_{k=1}^m e^{-\gamma^2 D_J(\mathbf{X}_i, \hat{\mathbf{X}}_k)}} = \frac{1}{\Omega_i} e^{-\gamma^2 D_J(\mathbf{X}_i, \hat{\mathbf{X}}_j)}, \quad (4)$$

where $D_J(\mathbf{X}_i, \hat{\mathbf{X}}_j)$ is the JBLD divergence in eq. (1) and Ω_i denotes the normalization term. The constant $\gamma > 0$ is a hyper-parameter defining the "sharpness" of the neighborhood distribution. (We set γ by cross-validation)

Objective. Inspired by Neighborhood Components Analysis [16], we can compute the probability p_i that an input $\mathbf{X}_i$ will be classified correctly by this stochastic nearest neighbor classifier under the compressed set $\{\hat{\mathbf{X}}_1, \ldots, \hat{\mathbf{X}}_m\}$,

$$p_i = \sum_{j:y_j = y_i} p_{ij}. \quad (5)$$

Ideally, $p_i = 1$ for all $\mathbf{X}_i$, implying the compressed set yields perfect predictions on the training set. The KL-divergence between this ideal "1-distribution" and p_i is simply $KL(1||p_i) = -\log(p_i)$. Our objective is to minimize the sum of these KL-divergences with respect to our compressed set of covariance matrices $\{\hat{\mathbf{X}}_1, \ldots, \hat{\mathbf{X}}_m\}$,

$$\min_{\{\hat{\mathbf{X}}_1, \ldots, \hat{\mathbf{X}}_m\}} -\sum_{i=1}^n \log(p_i). \quad (6)$$

Gradient. To ensure that the learned matrices $\hat{\mathbf{X}}_j$ are SPD, we decompose each matrix $\hat{\mathbf{X}}_j$ by its unique Cholesky decomposition: $\hat{\mathbf{X}}_j = \mathbf{B}_j^\top \mathbf{B}_j$, where $\mathbf{B}_j$ is an upper triangular matrix. To ensure that $\hat{\mathbf{X}}_j$ remains SPD we perform gradient descent w.r.t. $\mathbf{B}_j$. The gradient of $\mathcal{L}$ w.r.t. $\mathbf{B}_j$ is

$$\frac{\partial \mathcal{L}}{\partial \mathbf{B}_j} = \sum_{i=1}^n \frac{p_{ij}}{p_i}(\delta_{y_i y_j} - p_i)\gamma^2 \frac{\partial D_J(\mathbf{X}_i, \mathbf{B}_j^\top \mathbf{B}_j)}{\partial \mathbf{B}_j} \quad (7)$$

where $\delta_{y_i y_j} = 1$ if $y_i = y_j$ and is 0 otherwise and $D_J(\mathbf{X}_i, \mathbf{B}_j^\top \mathbf{B}_j)$ is the JBLD divergence between $\mathbf{X}_i$ and $\mathbf{B}_j^\top \mathbf{B}_j = \hat{\mathbf{X}}_j$. The gradient of the JBLD w.r.t. $\mathbf{B}_j$ is:

$$\frac{\partial D_J(\mathbf{X}_i, \mathbf{B}_j^\top \mathbf{B}_j)}{\partial \mathbf{B}_j} = \mathbf{B}_j(\mathbf{X}_i + \mathbf{B}_j^\top \mathbf{B}_j)^{-1} - (\mathbf{B}_j^\top)^{-1}. \quad (8)$$

We substitute (8) into (7) to obtain the final gradient. For a single compressed input $\hat{\mathbf{X}}_j = \mathbf{B}_j^\top \mathbf{B}_j$, each step of gradient descent requires $\mathrm{O}\big(d^3\big)$ to compute $\frac{\partial D_J(\mathbf{X}_i, \mathbf{B}_j^\top \mathbf{B}_j)}{\partial \mathbf{B}_j}$ and $\mathrm{O}\big(d^{2.37}\big)$ to compute $D_J(\mathbf{X}_i, \hat{\mathbf{X}}_j)$. It requires $\mathrm{O}\big(d^3 m\big)$ to compute p_{ij} and an additional $\mathrm{O}\big(m\big)$ for p_i. Thus the overall complexity of $\frac{\partial \mathcal{L}}{\partial \hat{\mathbf{X}}_j}$ is $\mathrm{O}\big(d^3 m^2 n\big)$. We minimize our objective in eq. (6) via conjugate gradient descent.[1] A Matlab implementation of SCC is available at: `http://anonymized`.

4. Histogram compression

Analogous to covariance compression, we can also compress histogram descriptors, which we refer to as *Stochastic Histogram Compression* (SHC). Our aim is to learn a compressed set of $m \ll n$ histograms $\{\hat{\mathbf{h}}_1, \ldots, \hat{\mathbf{h}}_m\} \subset \mathbf{\Sigma}_d$ with labels $\hat{y}_1, \ldots, \hat{y}_m$ from a training set histograms

[1] `http://tinyurl.com/minimize-m`

$\{\mathbf{h}_1, \ldots, \mathbf{h}_n\} \subset \Sigma_d$ with labels $y_1, \ldots, y_n$, where $\mathbf{\Sigma}_d$ is the $(d-1)$-dimensional simplex.

Objective. As before we place a stochastic neighborhood distribution over compressed histograms and define the probability that $\hat{\mathbf{h}}_j$ is the nearest neighbor of $\mathbf{h}_i$ via

$$p_{ij} = \frac{1}{\Omega_i} e^{-\gamma^2 D_S(\mathbf{h}_i, \hat{\mathbf{h}}_j)} \tag{9}$$

where D_S is the Sinkhorn distance and Ω_i normalizes p_{ij} so that it is a valid probability. As in SCC, we define the probability that a training histogram $\mathbf{h}_i$ is predicted correctly as p_i via eq. (5) by summing over the compressed inputs $\hat{\mathbf{h}}_k$ the share the same label. We then minimize the KL-divergence between the perfect distribution and p_i as in eq. (6) to learn our set of compressed histograms $\hat{\mathbf{h}}_j$.

Gradient. As in the covariance setting, the gradient of the objective in eq. (6) w.r.t. a compressed histogram $\hat{\mathbf{h}}_j$ is

$$\frac{\partial \mathcal{L}}{\partial \hat{\mathbf{h}}_j} = \sum_{i=1}^{n} \frac{p_{ij}}{p_i} (\delta_{y_i y_j} - p_i) \gamma^2 \frac{\partial D_S(\mathbf{h}_i, \hat{\mathbf{h}}_j)}{\partial \hat{\mathbf{h}}_j}. \tag{10}$$

The gradient of the Sinkhorn distance D_S w.r.t. $\hat{\mathbf{h}}_j$ introduces two challenges: 1. the distance D_S itself is a nested optimization problem; and 2. the learned vector $\hat{\mathbf{h}}_j$ must remain a well-defined histogram throughout the optimization, *i.e.* it must be non-negative and sum to 1, s.t. $\hat{\mathbf{h}}_j \in \mathbf{\Sigma}_d$.

We first address the gradient of the nested optimization problem w.r.t. $\hat{\mathbf{h}}_j$, *i.e.* $\frac{\partial D_S(\mathbf{h}_i, \hat{\mathbf{h}}_j)}{\partial \hat{\mathbf{h}}_j}$. In the *primal* formulation, as stated in eq. (3), the histogram $\hat{\mathbf{h}}_j$ occurs within the constraints, which complicates the gradient computation. Instead, we form the *dual* [8],

$$\max_{\boldsymbol{\alpha}, \boldsymbol{\beta} \in \mathbb{R}^d} \boldsymbol{\alpha}^\top \mathbf{h}_i + \boldsymbol{\beta}^\top \hat{\mathbf{h}}_j - \sum_{k,l=1}^{d} \frac{e^{-\lambda(M_{ij} - \alpha_i - \beta_j)}}{\lambda}, \tag{11}$$

where $\boldsymbol{\alpha}, \boldsymbol{\beta}$ are the corresponding dual variables. Due to strong duality, the primal and dual formulations are identical at the optimum, however the dual formulation (11) is *unconstrained*. The gradient of the dual objective (11) is *linear* w.r.t. $\hat{\mathbf{h}}_j$. If we consider $\boldsymbol{\beta}$ fixed, it follows that at the optimum $\frac{\partial D_S(\mathbf{h}_i, \hat{\mathbf{h}}_j)}{\partial \hat{\mathbf{h}}_j} \approx \boldsymbol{\beta}^*$, where $\boldsymbol{\beta}^*$ the optimal value of $\boldsymbol{\beta}$. This optimal dual variable is easily computed with the iterative Sinkhorn algorithm [9] mentioned in Section 2. This approximation ignores that $\boldsymbol{\beta}^*$ itself is a function of $\hat{\mathbf{h}}_j$, which is a reasonable approximation for small step-sizes.

To address the second problem and simultaneously perform (approximated) gradient descent while ensuring that $\hat{\mathbf{h}}_j$ always lies in the simplex (i.e., is normalized), we propose a change of variable in which we redefine each compressed histogram $\hat{\mathbf{h}}_j$ as a positive, normalized quantity: $\hat{\mathbf{h}}_j = e^{\mathbf{w}_j} / \sum_{k=1}^{d} e^{\mathbf{w}_{jk}}$ for $\mathbf{w}_j \in \mathbb{R}^d$. Then the gradient of the Sinkhorn distance can be taken with respect to $\mathbf{w}_j$,

$$\frac{\partial D_S(\mathbf{h}_i, \hat{\mathbf{h}}_j)}{\partial \mathbf{w}_j} \approx \boldsymbol{\beta}^* \circ \left(\frac{s e^{\mathbf{w}_j} - e^{2\mathbf{w}_j}}{s^2} \right)$$

where s is the normalizing term $s = \sum_{k=1}^{d} e^{\mathbf{w}_{jk}}$ and $\circ$ is the Hadamard (element-wise) product. The complexity of computing the full SHC gradient in eq. (10) is $\mathrm{O}\big(nmd^3\hat{i}\big)$: each Sinkhorn gradient above requires time $\mathrm{O}\big(d^3\hat{i}\big)$ (where $\hat{i}$ is the number of Sinkhorn iterations), and p_{ij} requires time $\mathrm{O}\big(\hat{i}d^2m\big)$. We use gradient descent with an updating learning rate to learn $\hat{\mathbf{h}}_j$, selecting the compressed set that yields the best training error across all iterations.

5. Results

We evaluate both algorithms on a series of real-world data sets and compare with state-of-the-art algorithms for kNN compression.

5.1. Covariance compression

Datasets. We evaluate our covariance compression method on six benchmark data sets. The ETH80 dataset has images of 8 object categories, each pictured with a solid blue background. For each category there are 10 exemplar objects and for each exemplar the camera is placed in 41 different positions. We use the 19×19 covariance descriptors of [5], who segment the images and use per-pixel color and texture features to construct covariances. The ETHZ dataset is a low-resolution set of images from surveillance cameras of sizes from 78×30 to 400×200. The original task is to identify the person in a given image, from 146 different individuals. The original dataset has multiple classes with fewer than 10 individuals. Therefore, to better demonstrate a wide range of compression ratios, we filter the dataset to include only the most popular 50 classes resulting in each individual having between 59 and 356 images (5193 images total). We use the pixel-wise features of [5] to construct 18×18 covariance matrices. The FERET face recognition dataset has 3737 gray-scale images of the faces of 399 individuals, oriented at various angles. As the majority of individuals have fewer than 5 images in the training set, we also limit the dataset to the 50 most popular individuals and use a larger set of compression ratios (described further in the error analysis subsection). We use the 40×40 Gabor-filter covariances of [6]. Our version of the RGBD Object dataset [25] contains $15,000$ point cloud frames of objects from three different views. The task is to classify an object as one of 51 object categories. We use the 18×18 covariance features of [5] which consist of intensity and depth-map gradients, as well as surface normal information. The SCENE15 data

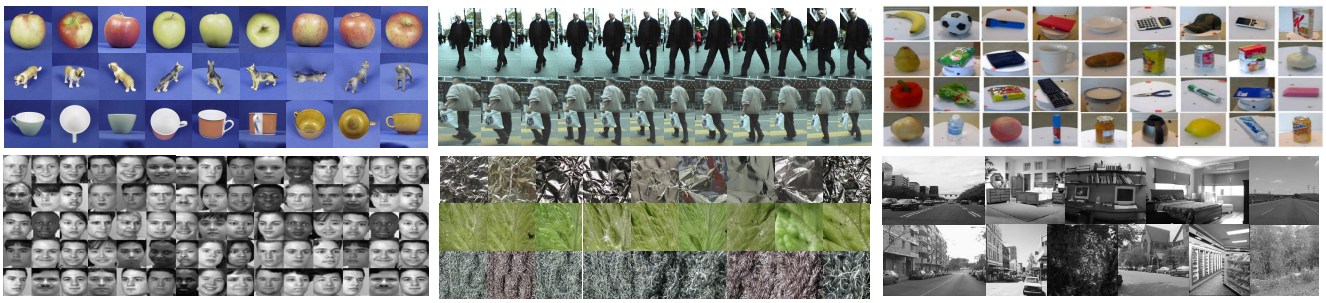

Figure 1. Montages of the 6 datasets used in our evaluations from top down, left to right: (a) ETH80 objects at different orientations; (b) ETHZ person recognition; (c) RGBD objects from point clouds; (d) FERET face detection; (e) KTH-TIPS2B material categorization (f) SCENE15 scene classification.

consists of 4485 black and white images of 15 different indoor and outdoor scenes. We split the dataset into training and test sets as per [28]. To create covariance features we compute a dense set of SIFT descriptors centered at each pixel in the image[2]. Our SIFT features have 4 bins each in the horizontal and vertical directions and 8 orientation bins, producing a 128×128 covariance descriptor. Via the work of [19] we can learn a rank-r projection matrix $\mathbf{U} \in \mathbb{R}^{d \times r}$, where $r \ll d$ to reduce the size of these covariance matrices to $r \times r$ via the transformation $\mathbf{X}_i \rightarrow \mathbf{U}^\top \mathbf{X}_i \mathbf{U}$. We use Bayesian optimization [13][3], to select values for the covariance size r, as well as two hyperparameters in [19]: ν_v and ν_w, by minimizing the 1-NN error on a small validation set. The KTH-TIPS2B dataset is a material classification dataset of 11 materials with 4752 total images. Each material has 4 different samples each from 3 fixed poses, 9 scales, and 4 lighting conditions. We follow the procedure of [18] to extract 23×23 covariance descriptors using color information and Gabor filter responses.

Experimental Setup. We compare all methods against test error of 1-nearest neighbor classification that uses the entire training set. For results that depend on random initialization or sampling we report the average and standard deviation across 5 random runs (save KTH-TIPS2B, for which we use 4 splits by holding out each of the four provided samples, one at a time). For datasets RGBD, ETHZ, and ETH80 we report results averaged over 5 different train/test splits.

Baselines. We compare our method, *Stochastic Covariance Compression* (SCC), against a number of methods aimed at reducing the size of the training set, which we adapt for the covariance feature setting: 1. *kNN* using the full training set, 2. kNN using a class-based *subsampled* training set, which we use as initialization for SCC, 3. *Condensed Nearest Neighbor* (CNN) [20], 4. *Reduced Nearest Neighbor* (RNN) [14], 5. *Random Mutation Hill Climbing* (RMHC) [42], and 6. *Fast CNN* (FCNN) [1]. Both CNN and FCNN select subsets of the training set that have the same leave-one-out training error as the full training set, and are very well-known in the fast kNN literature. RNN works by post-processing the output of CNN to improve the training error and RMHC is a random subset selection method. (We give further details on these algorithms in Section 6). For FCNN we must make a modification to accommodate covariance matrix features. Specifically, FCNN requires computing the centroid of each class at regular intervals during the selection. A centroid of class y is given by solving the following optimization,

$$\mathbf{X}_y = \arg\min_{\mathbf{X}} \sum_{i:y_i=y} D_J(\mathbf{X}, \mathbf{X}_i)$$

where $D_J(\mathbf{X}, \mathbf{X}_i)$ is the JBLD divergence. Cherian et al. [6] give an efficient iterative procedure for solving the above optimization, which we use in our covariance FCNN implementation.

Classification error. Figure 2 compares the test error of SCC to baselines for sizes of the compressed set equal to $2\%, 4\%, 8\%$, and 16% of the training set. For dataset FERET which has a large number of classes we use larger compression ratios: $10\%, 20\%, 30\%$, and 40%. Although CNN, RNN, and FCNN only output a single reduced training set (the final point on each curve) we plot the intermediate test errors of each method at the above compression ratios as well. On each dataset SCC is able to reduce the test error to nearly that of kNN applied to the full dataset using less than or equal to 20% of the training data. Only on ETHZ and RGBD could SCC not match the full kNN error up to significance, however the error rates are only marginally higher. For small compression ratios SCC is superior to all of the baselines, as well as the subsampling initialization. On datasets ETH80 and ETHZ the final outputs of CNN, RNN and FCNN are roughly equivalent to the SCC curve. However, one notable downside is that these algorithms have no control on the size of these final sets, which for FERET are as large as 75% (CNN/RNN) and 74% (FCNN). In contrast SCC allows one to regulate the compressed set size precisely. RMHC is also able to regulate the compressed set

[2]We use the open source library VLFeat http://www.vlfeat.org/

[3]https://bitbucket.org/mlcircus/bayesopt.m

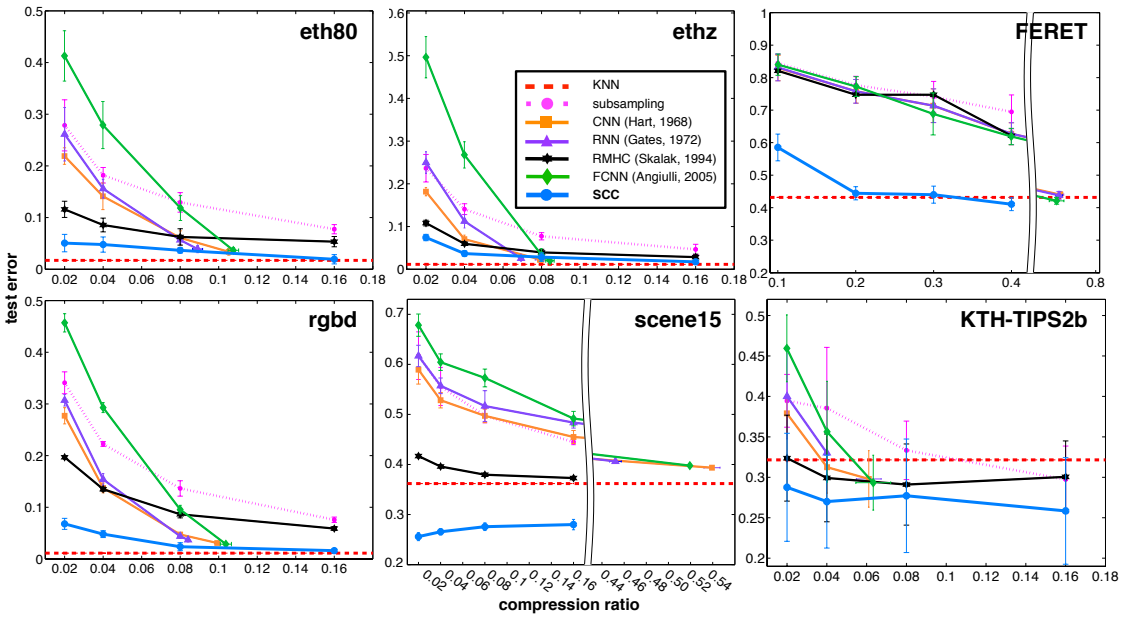

Figure 2. kNN test error rates after training set compression for covariance descriptors. See text for details.

size. However, because it is based on random sampling, its performance can be very poor, as on FERET. Surprisingly, for SCENE15 and KTH, learning a compressed covariance dataset with SCC reduces the test error *below the kNN error of the full training set*. We suspect that this occurs because (a) the training set may have some amount of label noise, which is partially alleviated by subsampling, and (b) SCC essentially learns a new, supervised covariance representation, versus a label-agnostic set of covariance descriptors. For these datasets, there is no reason not to shrink the dataset to only 2% of its original size and discard the original data—yielding $36\times$ and $48\times$ speedups during test-time.

Test-time speedup. Table 1 shows the speedup of SCC over kNN classification using the full training set for various compression ratios (the datasets marked by a (C) are learned with SCC). In general the speedups are roughly $1/\delta$, where δ denotes the compression ratio. Results that match or exceed the accuracy (up to significance) are in blue. At 16% compression 3 of the 5 datasets run at this compression ratio match the test error of full kNN classification. In effect we have removed neighbor redundancies in the dataset, and gained a factor of roughly $6\times$ speedup. Much larger speedups can be obtained at 4% or 2% compression ratio—although at a small increase in classification error. For the data set with many classes (FERET) "loss-free" compression can still yield a speedup of $5\times$ at $\delta = 0.2$.

Training time. Table 2 describes the average training times for SCC (again (C) denotes SCC results). For maximum compression to 2% the training time is on the order of minutes. As the size of the compressed set gets larger the time increases but only by small amounts, indeed the longest training time is 2 hours for RGBD with 16% compression. Furthermore, the entire compression can be done completely off-line prior to testing. The contributions of the training points to the gradient are independent and have a high computation to memory load ratio. The SCC training could therefore potentially be sped up significantly through parallelization on clusters or GPUs.

Table 1. Speed-up of kNN testing through SCC and SHC compression. The SCC datasets are denoted with a (C) and the SHC datasets with an (H). Results where SCC/SHC matches or exceeds the accuracy of full kNN (up to statistical significance) are in blue.

SPEED-UP				
DATASET	COMPRESSION RATIO			
(FEW CLASSES)	2%	4%	8%	16%
ETH80 (C)	47.6 ± 0.1	24.5 ± 0.1	12.3 ± 0.2	6.2 ± 0.02
ETHZ (C)	48.1 ± 1.0	24.5 ± 0.1	12.4 ± 0.05	6.2 ± 0.01
RGBD (C)	34.0 ± 0.8	20.3 ± 0.2	10.9 ± 0.2	6.8 ± 2.5
SCENE15 (C)	36.0 ± 0.6	20.6 ± 0.3	11.3 ± 0.1	5.9 ± 0.08
KTH (C)	43.0 ± 1.0	22.6 ± 0.7	12.0 ± 0.4	6.2 ± 0.2
COIL20 (H)	46.8 ± 1.4	23.5 ± 0.5	11.8 ± 0.5	5.9 ± 0.2
KYLBERG (H)	46.9 ± 2.2	23.9 ± 0.3	12.0 ± 0.04	6.0 ± 0.04
DATASET	COMPRESSION RATIO			
(MANY CLASSES)	10%	20%	30%	40%
FERET (C)	9.9 ± 0.1	5.0 ± 0.1	3.3 ± 0.006	2.5 ± 0.01
MPEG7 (H)	9.9 ± 0.2	4.9 ± 0.06	3.2 ± 0.01	2.4 ± 0.04

5.2. Histogram compression

We evaluate our technique for compressing histogram datasets, Stochastic Histogram Compression (SHC) against current baseline methods for constructing a reduced training set. As a benchmark, we compare the k-nearest neighbor accuracies for compressed sets of different sizes, and report the test-time speedups achieved by our method. We start by describing the datasets we use for comparison.

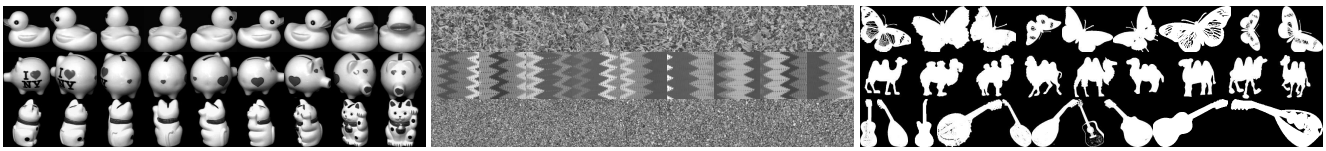

Figure 3. Montages of datasets used in SHC evaluation from left to right: (a) COIL20 3D object recognition; (b) KYLBERG texture classification; (c) MPEG7 shape detection.

Table 2. SCC and SHC training times.

TRAINING TIMES				
DATASET	COMPRESSION RATIO			
(few classes)	2%	4%	8%	16%
ETH80 (C)	1m 9s	1m 46s	2m 39s	4m 36s
ETHZ (C)	2m 4s	3m 15s	5m 53s	10m 10s
RGBD (C)	18m 26s	33m 30s	1h 4m	2h 11m
SCENE15 (C)	1m 18s	1m 52s	2m 42s	4m 20s
KTH (C)	3m 31s	5m 48s	9m 2s	16m 18s
COIL20 (H)	15s	3m 21s	4m 43s	9m 17s
KYLBERG (H)	1m 41s	2m 46s	5m 24s	11m 17s
DATASET	COMPRESSION RATIO			
(many classes)	10%	20%	30%	40%
FERET (C)	1m 39s	2m 19s	2m 44s	3m 6s
MPEG7 (H)	58s	19m 26s	29m 33s	38m 52s

Datasets. The COIL20 dataset consists of 20 grayscale image objects with background masked out in black. Each object was rotated 360 degrees and an image was taken every 5 degrees, yielding 72 images per class. To construct histogram features we follow the procedure of [2] to extract shape context log-polar histograms using 100 randomly sampled edge points, yielding histograms of dimensionality $d = 60$. As a ground distance $\mathbf{M}$ we use the ℓ_1 distance between bins of the log-polar histogram. The MPEG7 dataset has 70 different shape classes, each with 20 images. Each image has a black background with a sold white shape such as *bat*, *cellular phone*, *fountain*, and *octopus*, among others. We follow the procedure for the COIL20 dataset to extract shape context histograms, also used in [2] for the MPEG7 dataset. The ground distance is also the ℓ_1 between bins. The KYLBERG texture dataset is a 28-class dataset of different surfaces. We used the dataset without rotations that contains 160 images for each class. We follow the feature-extraction technique of [18], which uses first and second order image-gradient features at every 4 pixels, after resizing. We then construct a visual bag-of-words representation by first clustering all features into 50 codewords. We represent each image as a 50-dimensional count vector: the i^{th} entry corresponds to the number of times a gradient feature was closest (in the Euclidean sense) to the i^{th} codeword. As a ground distance between bins we use the Euclidean distance between each pair of codewords.

Experimental setup. As for covariance features, our benchmark for comparison of all methods is the test error of 1-nearest neighbor classification with the full training set. Similarly, for each dataset we report results over 5 different train/test splits. For our algorithm, SHC, we use Bayesian optimization [13] to tune the γ^2 parameter in the definition of p_{ij}, eq. (9), as well as the initial gradient descent learning rate, to minimize the training error. Additionally, we initialize SHC with the results of RMHC, which in the covariance setting appears to largely outperform the subsampling approach. We use the exact same baselines for covariance features, except now we use the Sinkhorn distance as our dissimilarity measure. The only subtlety is that FCNN needs to be able to compute the centroid of a set of histograms, with respect to the Sinkhorn distance. The centroid of a set of histogram measures with respect to the EMD is called the *Wasserstein Barycenter* [9]. It is shown how to compute this barycenter for the Sinkhorn distance in [8], and we use their accelerated gradient approach to solve for each Sinkhorn centroid.

Classification error. Figure 4 shows the average test error and standard deviation for compression ratios of $2\%, 4\%, 8\%$, and 16% for COIL20 and KYLBERG and $10\%, 20\%, 30\%, 40\%$ for the many-class dataset MPEG7. for each of the above datasets. As in the covariance setting, SHC outperforms or matches the error vs. compression trade-offs of all of the baseline methods throughout all evaluated settings. For a specific speedup over full kNN classification (*i.e.* compression ratio), SHC is able to achieve the lowest test error (possibly matched by other methods) throughout. On KYLBERG, SHC can reduce the training set to $1/50$ of its size without an increase in test error. On COIL20 and MPEG7 the final compressed sets of CNN, RNN and FCNN have very high compression ratio (around 0.3 and 0.5), which lead to only very modest speedups. We did not evaluate SHC in these arguably least interesting settings, but nevertheless show the error rates of the baselines for completeness.

Test-time speedup. The kNN test time speedups of SHC over the full training set are shown in Table 1 (the (H) datasets). Similar to SCC, the speedups reach up to $46.9\times$ at maximum compression and still reach an order of magnitude in the worst case ($10\times$) on the MPEG7 data set with many classes. For KYLBERG the SHC error is lower than the full data set—even at a 2% compression ratio.

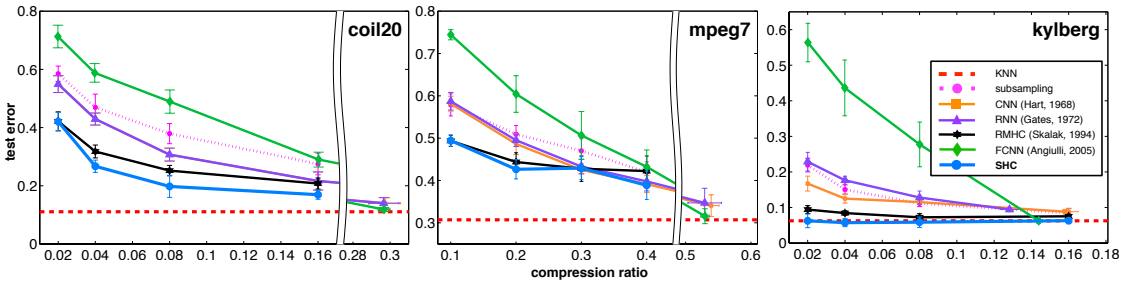

Figure 4. kNN test error rates after training set compression for histogram descriptors. See text for details.

Training time. The training times of SHC are shown in Table 2. SHC is very fast (< 2 minutes for 2% compression on KYLBERG), especially considering that we are solving a nested optimization problem over the compressed histograms and the Sinkhorn distance. We believe that the speed of our implementation can be further improved with the use of GPUs for the Sinkhorn computation and with approximate second-order hill-climbing methods.

6. Related Work

Training set reduction has been considered in the context of kNN for vector data and the Euclidian distance with three primary methods: (a) *training consistent sampling*, (b) *prototype generation* and (c) *prototype positioning* (for a survey see [44]). Training consistent sampling iteratively adds inputs from the training set to a reduced 'reference set' until the reference set is perfectly classified by the training set. This is precisely the technique of Condensed Nearest Neighbors (CNN) [20]. There have been a number of extensions of CNN, notably Reduced Nearest Neighbor (RNN) [14] which searches for the smallest subset of the result of CNN that correctly classifies the training data. Additionally, Fast CNN (FCNN) [1] finds a set close to that of CNN but has training time linear (instead of cubic) in the size of the training set. Prototype generation creates new inputs to represent the training set, usually via clustering [23, 40]. Prototype positioning learns a reduced training set by optimizing an appropriate objective. The method most similar to SCC and SHC is the recently proposed Stochastic Neighbor Compression [24], which uses a stochastic neighborhood to learn prototypes in Euclidean space (and thus is unsuitable for covariance and histogram features). Finally, Bucilua et al. [3] may have been the first to study model compression for machine learning algorithms by compressing neural networks. To our knowledge, SCC and SHC are the first methods to explicitly consider training set reduction for SPD covariance and histogram descriptors.

Work towards speeding up test-time classification for kNN on covariance-valued data is somewhat limited. The JBLD divergence is proposed to speed up individual distance computations. Cherian et al. [6] show that it is possible to adapt Bregman Ball Trees (BBTs), a generalization of the Euclidean ball tree to Bregman divergences, to the JBLD divergence. This is done using a clever iterative K-means method followed by a leaf node projection technique onto relevant Bregman balls. Both of these techniques are complementary to our dataset compression method.

There has been a large amount of work devoted toward improving the complexity of the Earth Mover's distance using approximations [8, 17, 41, 30, 31, 35]. For instance, [35] point out that if an upper bound can be placed on the transport $\mathbf{T}_{ij}$ between any two bins i and j, then the (thresholded) EMD can be solved much more efficiently. Ling and Okada [31] show that if the ground distance is the ℓ_1 distance between bins, the EMD can be reformulated exactly as a tree-based optimization problem with d unknown variables (instead of d^2) and only d constraints (instead of $2d$). We use the Sinkhorn approximation [8], which has the added advantage of an unconstrained dual formulation.

7. Conclusion

In many classification settings the sheer amount of distance computations has previously prohibited the use of kNN for covariance and histogram features. We have shown that these data sets can be compressed to a small fraction of their original sizes while often only slightly increasing the test error. This drastically speeds up nearest neighbor search and has the potential to unlock new applications for covariance and histogram features on large datasets.

References

[1] F. Angiulli. Fast condensed nearest neighbor rule. In *ICML*, pages 25–32, 2005.

[2] S. Belongie, J. Malik, and J. Puzicha. Shape matching and object recognition using shape contexts. *TPAMI*, 24(4):509–522, 2002.

[3] C. Bucilua, R. Caruana, and A. Niculescu-Mizil. Model compression. In *SIGKDD*, pages 535–541. ACM, 2006.

[4] L. Cayton. Fast nearest neighbor retrieval for Bregman divergences. In *ICML*, pages 112–119. ACM, 2008.

[5] A. Cherian and S. Sra. Riemannian sparse coding for positive definite matrices. In *ECCV*, pages 299–314. Springer, 2014.

[6] A. Cherian, S. Sra, A. Banerjee, and N. Papanikolopoulos. Efficient similarity search for covariance matrices via the

Jensen-Bregman LogDet divergence. In *ICCV*, pages 2399–2406, 2011.

[7] D. Coppersmith and S. Winograd. Matrix multiplication via arithmetic progressions. *Journal of symbolic computation*, 9(3):251–280, 1990.

[8] M. Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In *NIPS*, pages 2292–2300, 2013.

[9] M. Cuturi and A. Doucet. Fast computation of Wasserstein barycenters. In *ICML*, 2014.

[10] N. Dalal and B. Triggs. Histograms of oriented gradients for human detection. In *CVPR*, volume 1, pages 886–893, 2005.

[11] L. Fei-Fei, R. Fergus, and A. Torralba. Recognizing and learning object categories. *CVPR Short Course*, 2, 2007.

[12] L. Fei-Fei and P. Perona. A bayesian hierarchical model for learning natural scene categories. In *CVPR*, volume 2, pages 524–531, 2005.

[13] J. Gardner, M. Kusner, Z. Xu, K. Q. Weinberger, and J. Cunningham. Bayesian optimization with inequality constraints. In *ICML*, pages 937–945, 2014.

[14] G. Gates. The reduced nearest neighbor rule. *IEEE Transactions on Information Theory*, 18:431–433, 1972.

[15] A. Goh and R. Vidal. Clustering and dimensionality reduction on Riemannian manifolds. In *CVPR*, pages 1–7, 2008.

[16] J. Goldberger, G. Hinton, S. Roweis, and R. Salakhutdinov. Neighbourhood components analysis. In *NIPS*, pages 513–520. 2004.

[17] K. Grauman and T. Darrell. Fast contour matching using approximate earth mover's distance. In *CVPR*, volume 1, pages I–220, 2004.

[18] M. Harandi, M. Salzmann, and F. Porikli. Bregman divergences for infinite dimensional covariance matrices. In *CVPR*, pages 1003–1010, 2014.

[19] M. T. Harandi, M. Salzmann, and R. Hartley. From manifold to manifold: geometry-aware dimensionality reduction for SPD matrices. In *ECCV*, pages 17–32. 2014.

[20] P. Hart. The condensed nearest neighbor rule. *IEEE Transactions on Information Theory*, 14:515–516, 1968.

[21] G. Hinton and S. Roweis. Stochastic neighbor embedding. In *NIPS*, pages 833–840. 2002.

[22] S. Jayasumana, R. Hartley, M. Salzmann, H. Li, and M. Harandi. Kernel methods on the Riemannian manifold of symmetric positive definite matrices. In *CVPR*, pages 73–80, 2013.

[23] T. Kohonen. Improved versions of learning vector quantization. In *IJCNN*, pages 545–550. IEEE, 1990.

[24] M. Kusner, S. Tyree, K. Q. Weinberger, and K. Agrawal. Stochastic neighbor compression. In *ICML*, pages 622–630, 2014.

[25] K. Lai, L. Bo, X. Ren, and D. Fox. A large-scale hierarchical multi-view rgb-d object dataset. In *ICRA*, pages 1817–1824, 2011.

[26] K. I. Laws. Rapid texture identification. In *24th Annual Technical Symposium*, pages 376–381. International Society for Optics and Photonics, 1980.

[27] S. Lazebnik, C. Schmid, and J. Ponce. A sparse texture representation using local affine regions. *TPAMI*, pages 1265–1278, 2005.

[28] S. Lazebnik, C. Schmid, and J. Ponce. Beyond bags of features: Spatial pyramid matching for recognizing natural scene categories. In *CVPR*, pages 2169–2178, 2006.

[29] T. Leung and J. Malik. Representing and recognizing the visual appearance of materials using three-dimensional textons. *IJCV*, 43(1):29–44, 2001.

[30] E. Levina and P. Bickel. The earth mover's distance is the mallows distance: Some insights from statistics. In *ICCV*, pages 251–256, 2001.

[31] H. Ling and K. Okada. An efficient earth mover's distance algorithm for robust histogram comparison. *TPAMI*, 29(5):840–853, 2007.

[32] D. G. Lowe. Distinctive image features from scale-invariant keypoints. *IJCV*, 60(2):91–110, 2004.

[33] K. Mikolajczyk and C. Schmid. A performance evaluation of local descriptors. *TPAMI*, 27(10):1615–1630, 2005.

[34] C. W. Niblack, R. Barber, W. Equitz, M. D. Flickner, E. H. Glasman, D. Petkovic, P. Yanker, C. Faloutsos, and G. Taubin. Qbic project: querying images by content, using color, texture, and shape. In *IS&T/SPIE Symposium on Electronic Imaging*, pages 173–187, 1993.

[35] M. Pele, O.and Werman. Fast and robust earth mover's distances. In *ICCV*, pages 460–467. IEEE, 2009.

[36] O. Pele and M. Werman. The quadratic-chi histogram distance family. In *ECCV*, pages 749–762. Springer, 2010.

[37] X. Pennec, P. Fillard, and N. Ayache. A riemannian framework for tensor computing. *IJCV*, 66(1):41–66, 2006.

[38] Y. Rubner, J. Puzicha, C. Tomasi, and J. M. Buhmann. Empirical evaluation of dissimilarity measures for color and texture. *Computer vision and image understanding*, 84(1):25–43, 2001.

[39] Y. Rubner, C. Tomasi, and L. J. Guibas. The earth mover's distance as a metric for image retrieval. *IJCV*, 40(2):99–121, 2000.

[40] S. Salzberg, A. Delcher, D. Heath, and S. Kasif. Best-case results for nearest-neighbor learning. *PAMI*, 17(6):599–608, 1995.

[41] S. Shirdhonkar and D. W. Jacobs. Approximate earth mover's distance in linear time. In *CVPR*, pages 1–8, 2008.

[42] D. B. Skalak. Prototype and feature selection by sampling and random mutation hill climbing algorithms. In *ICML*, pages 293–301, 1994.

[43] M. A. Stricker and M. Orengo. Similarity of color images. In *IS&T/SPIE Symposium on Electronic Imaging*, pages 381–392, 1995.

[44] G. T. Toussaint. Proximity graphs for nearest neighbor decision rules: recent progress. *Interface*, 34, 2002.

[45] O. Tuzel, F. Porikli, and P. Meer. Region covariance: A fast descriptor for detection and classification. In *ECCV*, pages 589–600. Springer, 2006.

[46] L. Van der Maaten and G. Hinton. Visualizing data using t-sne. *JMLR*, 9(2579-2605):85, 2008.

[47] R. Vemulapalli and D. W. Jacobs. Riemannian metric learning for symmetric positive definite matrices. *arXiv preprint arXiv:1501.02393*, 2015.

[48] K. Weinberger and L. Saul. Distance metric learning for large margin nearest neighbor classification. *JMLR*, 10:207–244, 2009.